



---

# Extreme Scale Computing, Co-design

---

## Informing system design, ensuring productive and efficient code

---

### Project Description

To address the increasingly complex problems of the modern world, scientists at Los Alamos are pushing the scale of computing to the extreme, forming partnerships with other national laboratories and industry to develop supercomputers that can achieve “exaflop” speeds—that is, a quintillion (a million trillion) calculations per second. To put such speed in perspective, it is equivalent to 50 million laptops all working together at the same time. Researchers are also developing the interacting components of a computational system as a whole. This approach, known as computational co-design, may facilitate revolutionary designs in the next generation of supercomputers.

Since the late 1940s, scientists at Los Alamos have pushed the scale of computing to the extreme. To design the world’s first atomic bomb, scientists created the Monte Carlo method of computing. In 1948, work began on MANIAC, one of the first electronic, digital computers. In 2008, Los Alamos unveiled the Roadrunner supercomputer, which was the first to break the “petaflop” barrier.

Working with other national laboratories and industry, researchers are developing supercomputers that can perform a quintillion (a million trillion) calculations per second—known as the exascale. Such speed will enable scientists to address extremely large datasets and extremely high rate data streams needed to address problems in national security, cyber security, energy security, global climate modeling, astrophysics, and biology.

Computational co-design involves developing the interacting components of a computational system as a whole. Such an approach produces significantly better, perhaps even revolutionary designs in the next generation of extreme computers. Such work is primarily focused on how to co-optimize constraints related to physics, method, implementation, and architecture.

### Capabilities

- Extreme-Scale Computer Development and Operation
- Data Science at Scale
- Coupled Computational Physics Applications and Simulations at Scale
- Computational Co-design
- Complex Networks
- Next-Generation File Systems

- Advanced Digital Libraries

Technologies and Applications: Emerging, Developed, or Potential

## Accomplishments

- Deployed Luna, a supercomputer designed to support work for the Directed Stockpile Work Program, including the B61 Life-Extension Program. This supercomputer has a total of 24,640 processors for a combined peak capability of 539.1 teraflops/second. Such speeds mean that users can obtain faster turnarounds on their calculations and that the computer can run with higher fidelity, thus improving the end results. Such features are particularly advantageous when performing weapons safety calculations.
- Worked with Sandia National Laboratories under the Advanced Computing at Extreme Scale (ACES) partnership to design and develop the supercomputer [Cielo](#) (Spanish for “sky”), which was built by Cray Inc. Cielo can perform more than one quadrillion floating-point operations per second. Cielo supports work related to stockpile stewardship.
- Established in 2011 a Co-design Summer School that brought together a small, multidisciplinary team of students of various universities. These students focused on problems related to computational co-design. The goal of this summer school is to encourage qualified students to work in computational co-design as the world approaches the exascale era.
- Created CoCoMANS (Computational Co-design for Multiscale Applications in the Natural Sciences), a project designed to forge a qualitatively new predictive-science capability that exploits evolving high-performance computer architectures for multiple application areas, including materials, plasmas, and climate, by simultaneously evolving science, methods, software, and hardware in an integrated computational co-design process.
- Designed Cruft, a suite of molecular dynamics proxy applications (software) developed to explore co-design opportunities with hardware vendors and scientists. This code enables researchers to conduct different elements of a molecular dynamics application and explore the ramifications on relative performance when changing hardware and software. Cruft was developed to help guide design decisions for codes written to support the Exascale Co-design Center for Materials in Extreme Conditions.
- Collaborating with EMC Corporation to support the Department of Energy’s Exascale Initiative, which is aimed at boosting high-performance computing levels to the exaflops—a thousand times faster than current petascale capabilities. One effort includes the development of an open-source, extremely scalable data-management middleware library known as PLFS (Parallel Log Structured File System), which will be used on computing platforms that range from small clusters to the largest supercomputers in the world.
- Working to develop a framework for hardware-software co-design as a formally posed optimization problem. Although the optimization framework will apply to multiple problem domains, for the target application scientists will use molecular dynamics, an exemplar for the need for computational scaling. Scientists working under this project view co-design as search and selection from a vast space of hardware and software designs that map to performance metrics. The objective function designed for optimization has as main components run time

(or computational rate), problem size, simulated time duration, energy use, and hardware cost.

## LANL Facilities and Resources

- Strategic Computing Complex: Also known as the Nicholas C. Metropolis Center for Modeling and Simulation, this complex houses supercomputers that support the calculation, modeling, simulation, and visualization of complex nuclear weapons data in support of the Stockpile Stewardship Program. The complex includes a Data Visualization Corridor, which enables scientists to view the models and simulations created by the supercomputers. The Data Visualization Corridor includes a Powerwall Theater and a five-sided CAVE Immersive Laboratory, as well as desktop visualization and collaborator capabilities.

## Key Personnel at LANL

- Timothy Germann: Exascale Co-design Center for Materials in Extreme Environments
- Allen McPherson: Computer Science Lead for the Exascale Co-design Center for Materials in Extreme Environments

## Principal Contacts

- John Sarrao: Associate Director for Theory, Simulation, and Computation
- AnnMarie Cutler: Communications for Theory, Simulation, and Computation
- Susan Seestrom: Associate Director for Experimental Physical Sciences
- Karen Kippen: Communications for Experimental Physical Sciences
- Manuel Vigil: Trinity Project Director
- Gary Grider: ACES Co-director

## Sponsors, Funding Sources, or Agencies

- Department of Homeland Security
- Department of Defense
- Department of Energy Office of Advanced Scientific Computing Research
- Department of Energy Office of Science

## Publications

James Powell, Linn Collins, Ariane Eberhardt, David Izraelevitz, Jorge Roman, Thomas Dufresne, Mark Scott, Miriam Blake, and Gary Grider, "At scale" author name matching with Hadoop/MapReduce," *Library Hi Tech News* 29(4), 6–12 (2012).

Ning Liu, Jason Cope, Philip Carns, Christopher Carothers, Robert Ross, Gary Grider, Adam Crume, and Carlos Maltzahn, "On the role of burst buffers in leadership-class storage systems," *IEEE Symposium on Mass Storage Systems and Technologies* (2012).

John Bent, Gary Grider, Brett Kettering, Adam Manzanares, Meghan McClelland, Aaron Torres, and Alfred Torrez, "Storage challenges at Los Alamos National Lab," *IEEE Symposium on Mass Storage Systems and Technologies* (2012).

John Bent, Sorin Faibish, Jim Ahrens, Gary Grider, John Patchett, Percy Tzelnic, and Jon Woodring, "Jitter-free co-processing on a prototype exascale storage stack," *IEEE Symposium on Mass Storage Systems and Technologies* (2012).

Nithin Nakka, Alok Choudhary, Gary Grider, John Bent, James Nunez, and Satsangat Khalsa, "Achieving target MTTF by duplicating reliability-critical components in high performance computing systems," *IEEE International Symposium on Parallel and Distributed Processing Workshops and PhD Forum*, 1567–1576 (2011).

Matthew L. Curry, H. Lee Ward, Gary Grider, Jill Gemmill, Jay Harris, and David Martinez, "Power use of disk subsystems in supercomputers," *PDSW'11 - Proceedings of the 6th Parallel Data Storage Workshop, Co-located with SC'11*, 49–53 (2011).

Eugene Normand, Jerry L. Wert, Heather Quinn, Thomas D. Fairbanks, Sarah Michalak, Gary Grider, Paul Iwanchuk, John Morrison, Stephen Wender, and Steve Johnson, "First record of single-event upset on ground, Cray-1 computer at Los Alamos in 1976," *IEEE Transactions on Nuclear Science* 57(6 PART 1), 3114–3120 (2010).

Hsing-Bung Chen, Gary Grider, Cody Scott, Milton Turley, Aaron Torres, Kathy Sanchez, and John Bremer, "Integration experiences and performance studies of a COTS parallel archive system," *IEEE International Conference on Cluster Computing*, 166–177 (2010).

S.M. Mniszewski, M.J. Cawkwell, and T.C. Germann, "Molecular dynamics simulations of detonation on the Roadrunner supercomputer," *AIP Conference Proceedings* 1426, 1283–1286 (2012).

Christian Brandl and Timothy C. Germann, "Shock loading and release of a small angle tilt grain boundary in CU," *AIP Conference Proceedings* 1426, 1299–1302 (2012).

[More publications](#)

Los Alamos National Laboratory

[www.lanl.gov](http://www.lanl.gov)

(505) 667-7000

Los Alamos, NM

Operated by Los Alamos National Security, LLC for the Department of Energy's NNSA

